LA-UR -81-2587

CONF-811118--1

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

TITLE LONG-RANGE PREDICTION OF NETWORK TRAFFIC

AUTHOR(S) William Alexander
Richard Brice

DISCLAIMER

SUBMITTED TO Computer Performance Evaluation Users Group (CPEUG)

MASTER

By acceptance of this article the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes.
The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Los Alamos
Los Alamos National Laboratory
Los Alamos, New Mexico 87545

FORM NO. 836 R4
ST NO. 2629 5/81

LONG-RANGE PREDICTION OF NETWORK TRAFFIC

William Alexander
Richard Brice

Los Alamos National Laboratory
Computing Division
Los Alamos, NM 87544

A method of making long-range computer system workload predictions is presented. The method quantifies the effect of qualitative changes in computing by identifying assumptions and by considering the effect of a change on individual users. The method is illustrated by an example involving message traffic in a large computer network.

Key words: Computer networks; long-range forecasting; user behavior; workload forecasting.

1. Introduction

Management planning procedures sometimes equire computer system workload forecasts or five or even ten years in the future. resent workload prediction methods are nadequate at such long ranges because the hanges in system use are likely to be qualiative rather than just quantitative. erhaps a computer measurement professional hould be reluctant to make such long-range redictions if, possibly, too much credence ill be given them. However, when you must ake these predictions, how do you proceed?

In this paper we present a method for aking long-range workload predictions that uantifies the effects of qualitative changes n computing. Naturally, such predictions re somewhat speculative, and we claim only o provide a framework with which to organize nd quantify assumptions. The method consists of constructing a polynomial expression or the workload in which each term epresents the effects of one change. The erms are constructed by concentrating on the ffect that the change will have on individul users. This method explicitly represents ssumptions and allows parametric ranges of esults.

Some papers on workload forecasting for management planning look at current workload analysis, others study the extent of growth or change in computing activities. Determinat... of the current workload is heavily represented, probably because it is the most straightforward process in forecasting. Prediction methods include measurement techniques [1-5],[1] abstraction of synthetic workloads from the measurements [4,6,7], and reduction of the measurement data to manageable magnitude, for example, clustering analysis [8-10]. A second category addresses forecasting from a management perspective. These papers attempt to determine growth or change in activities that may affect computing needs. Isolated approaches exist that attempt to bridge the gulf between qualitative changes in the activities and their quantitative effects on computer resource requirements. Prediction methods in this area include extrapolation from resource requirements of existing application programs [11-13], forecasting of resource requirements for applications that are not yet completely implemented [14,15], and also some effects of

[1]Figures in brackets indicate the literature references at the end of this paper.

feedback between workload and level of service provided to users [16].

The essence of these approaches is to determine the nature of the current computing workload and, using this information, to project the amount of similar work that will be done at some future time. There are differences in how the current workload determinations are made and in the fundamental units of measure used to describe the workload. The units of measure range from resource utilization data for specific computer components to characterizations of project activities. These approaches seem best suited for short term (1-2 year) forecasts, because the effects of quantitative changes are likely to outweigh the effects of qualitative changes during this interval.

These methods are not suited to our specific problem, which is to forecast effects of qualitative changes in computing. In particular, these approaches do not address the influence that revolutions in computing hardware and services exert on how a user does his work. A second difficulty in forecasting is that long-term forecasts are almost certain to be wrong. This difficulty suggests that these forecasts should be cast in a form that is easy to update as new information arrives. Some problems that may occur if updating is not anticipated are described in Reference 17.

In Part II, we describe the need to predict message traffic in the Los Alamos National Laboratory (Los Alamos) Integrated Computing Network up to 1990. In Part III, we explain our method and illustrate its use.

2. The Problem

2.1 Integrated Computing Network

At Los Alamos, the Central Computing Facility (CCF) includes an Integrated Computing Network (ICN) that allows all validated computer users at the Laboratory access to almost any of the machines or services of the CCF.

Figure 1 is a schematic diagram of the ICN. At the "front end" of the Network (the right side of the diagram) an arbitrary

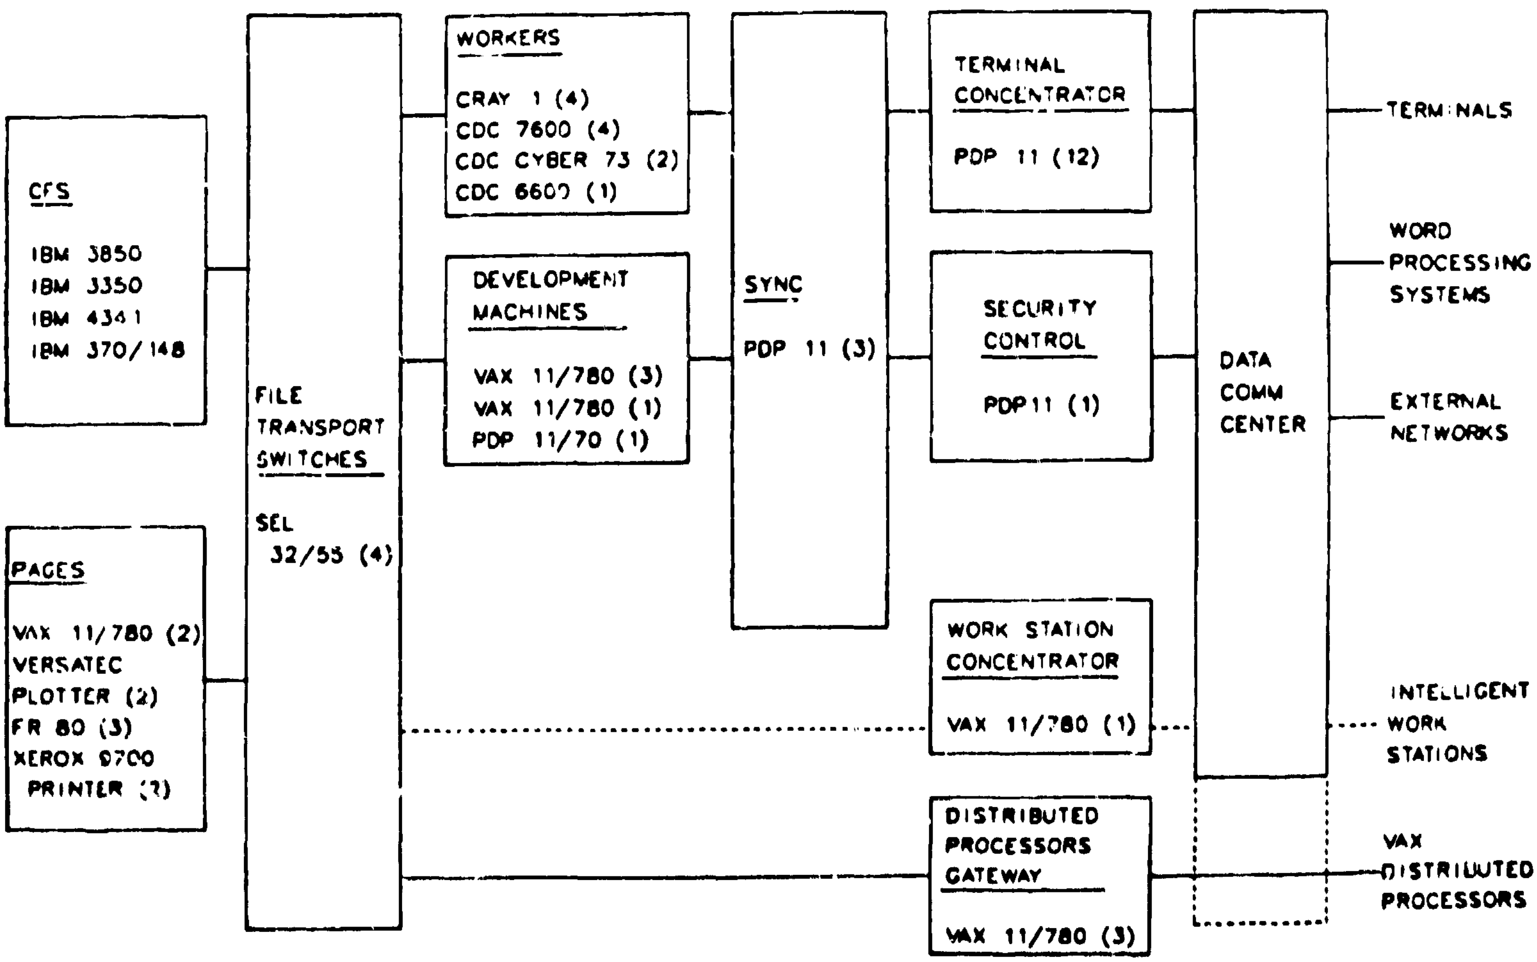

Figure 1. Functional Diagram of the ICN

number of terminals (currently about 1350) and remote entry stations are concentrated in stages to front end switches (the SYNCs), so that traffic can be routed between any terminal and any worker computer. Thus, aside from administrative restrictions, a user can log in on any worker from any terminal. The worker computers include three Cray-1s, four CDC 7600s, one CDC 6600, and two CDC Cyber-73 computers. Each of the worker computers is connected to the File Transport (FT) switches and, by the FT switches, to the "back end" of the Network (left side of the diagram). The FTs allow the workers to send files to each other and to the special service nodes in the Network. The special services provided by the Network at present include

- o an output station (PAGES) to which are attached a wide variety of printing and graphics devices,
- o a mass storage and archival facility (CFS) [18], and
- o XNET, which handles file traffic between workers and computers outside the ICN.

Messages between workers and SYNCs are usually quite small and are never larger than 1000 bytes. Messages routed through the FTs can be as large as 25,976 bytes; large files are broken by the sending machine into messages no larger than this, and the messages are sent sequentially (the ICN is not a packet-switching network).

In this paper, a "message" is one user- or program-defined group of bytes (plus network header) transferred together from a source node to a destination node in the ICN. "Nodes" include terminals, worker computers, and special service stations, but not concentrators or switches (SYNCs and FTs). From the point of view of network implementation, messages are certainly the appropriate unit of workload. From a larger view, considering the ICN as a unit, one might first think of workload in terms of terminal sessions, tasks submitted for execution on worker computers, etc. We believe that messages are an additional valid measure of workload, because there is a fairly direct correspondence between user or user program commands and messages generated. Messages result from a carriage return at the terminal and from certain explicit program functions or worker computer command language commands.

With colleagues we have just begun a new network performance measurement and evaluation project on the ICN. This project includes measurement and characterization of message traffic in the Network and analytic and simulation models of the Network. With these models we are beginning to identify the critical resources in the Network as well as to investigate the effects of increased traffic load, new equipment, and alternate configurations. Both the measurements and the models are at present rather crude.

In some of the following analysis we treat short messages (less than 100 bytes) and long messages separately, because our models indicate that different resources are critical in handling them. The critical resource limiting the Network's capacity to carry large messages seems to be buffer space in the switches, while line capacity and switch processor capacity are critical for small messages.

At present there are about 3000 users of the CCF. We measure approximately 20 large and 80 small messages per second in the back end of the Network and about 100 small messages per second in the front end. This is 0.06 small and 0.007 large messages per second per user.

2.2 The Forecasting Assignment

Recently we were asked by management to predict what the network traffic in the ICN would be at various points in the future up to the year 1990, 10 years from now. Current management forecasts indicate that the number of users of the Network will grow linearly from the present 3000 to 5700 in 1990; managers also anticipate a certain number of large worker computers in the Network by that year.

If the kind of work people do and how they go about doing it both remained constant, then the problem would be relatively straightforward. We might, for example, simply predict that the load in 1990 would be

$$(5700/3000) * (\text{present load})$$

ignoring the different number and kinds of worker machines in 1990 on the assumption that messages are generated by programs and people, not primarily by machines. However, computing habits have changed significantly in the past 10 years, and they are likely to again in the next 10. Timesharing radically altered the way people used computers in the 1970s, distributed processing and networks are doing it now, and there may be time for two more revolutions by 1990. Change seems to be a given in computing, and no one has developed a model to predict it. Thus we preceded our response with numerous caveats, and, when management promised to heed them, Clearly, the traditional PME predictive [...]l, namely a model of the Network, does not [...]ply directly to this problem; models are [...]signed to take workload as input, not to [...]edict it. Furthermore, there exists at [...]esent no characterization of our computer [...]rkload in terms of "worksteps" or "activity [...]its" [12,13], nor any formula for translat[...]g from these to network activity. Finally, [...]en if we had such workload characterization [...]d such a translation formula, it is not [...]ear that the formula would be valid for [...]mputing conditions 10 years hence. In [...]ct, the nature of the problem and the lack [...] data force us into the role of futurists [...] role [...]ich a systems analyst may be no [...]tter [...]alified than the next person.

3. The Solution

3.1 The Method

The central idea of our method is to [...]ncentrate on the individual user, that is, [...] predict the effect on the user of future [...]anges in network equipment, topology, and [...]rvices. This is clearly risky, because [...]eople are the least understood and least [...]redictable element in computing systems. [...]evertheless, this focus seems necessary, be[...]ause we do, in fact, believe that network [...]raffic is affected more by what people [...]hoose to do and how they choose to do it [...]han by the equipment they use. Of course, [...]etwork topology, equipment, and services [...]ake certain tasks easy and others more dif[...]icult, but so do other factors. We are not [...]rying to literally predict human behavior; [...]e are trying to orient and focus our think[...]ng in the face of too much uncertainty.

The first step is to identify factors [...]hat will change computing in our network. [...]hen we quantify the effect of each factor on [...]etwork traffic that individual users gen[...]rate. Finally, we collect the terms [...]epresenting each factor into a polynomial [...]xpression.

3.2 Five Factors

We were able to identify five factors [...]hat we believe will affect the way people [...]se the ICN in the next few years. They are [...]s follows.

1. Specialization of the Network. At present, CFS and PAGES are specialized nodes to which users from any worker can send files for permanent storage or for output. In the future, specialized nodes for word processing, for a network status and performance data base, and for other unanticipated functions may exist. (In fact, word processing software is available on a PDP-11/70 in the Network now, but this software is not yet widely used.) In addition, the worker computers themselves may become more specialized with some machines serving mostly as number crunchers and others as general-purpose front ends to the number crunchers.

2. Increased use of intelligent and graphics terminals.

3. Proliferation of distributed processors (DPs) and local networks of DPs within the Laboratory but outside the ICN. For a variety of reasons, the number of mini- and midicomputers outside the ICN continues to grow. They are used both for specialized purposes, such as process control, and for general computing; some are connected in small local networks. Typically these can communicate with any node in the ICN via XNET.

4. Electronic mail. Some electronic mail system will probably be installed at the Laboratory within the next few years, although it may be implemented as a separate mechanism rather than through the ICN.

5. Connections with remote networks. The most likely candidates are the computing facilities at other Department of Energy laboratories. Since these installations tend, at present, to have sufficient computing power for their own needs, the connections will probably be used to transfer data, programs, reports, etc., rather than to allow remote use of our computers. Similar connections to additional networks are possible.

Each of these five factors is either a trend that we see now in computing at Los Alamos or a capability currently being discussed and considered for inclusion here. In other words, we did not attempt any serious long-range crystal ball gazing, although the method allows this if you have the courage (see Section E). In the next section, we discuss the effect of each of these five factors on network message rates.

3.3 Analysis of Factors

It seems easiest to break the estimation of the effect that a change will have on any stem measure into two steps. First, one n analyze the qualitative aspects of the fect. For example, is the effect most narally expressed as a ratio to the present mber of messages a user generates or as an dition to that number? Is it independent the user's current activity? Is it inpendent of the number of users? Answers to ese questions will determine the position the factor, which represents a given ange in the polynomial formula for computg the value that the measure is expected to ve in the future. The second step is then plug in a numeric value for each factor, perhaps a range of numeric values.

We will illustrate this two-step process r each of the factors described in the preous section.

1. The specialization of the Network will clearly increase message rates. As specialized service nodes are added one by one, an individual user doing tasks functionally equivalent to present tasks will generate, perhaps even unknowingly, more network messages as his files are shipped to these nodes. The portion of a user's messages due to specialization will grow in proportion to the increased specialization of the Network. Therefore, a formula for the number of small messages in the Network should contain a multiplicative factor a in a term

   a*m*Ny,

   where Ny is the number of ICN users in the future year in question and m is the observed rate of small messages per user today. That is, specialization will increase small messages per user per unit time by some factor a. There will be a similar term

   A*M*Ny

   in the formula for large messages. The way specialized nodes are now used indicates that the users will mostly ship large files that will appear as large messages; this is partly a matter of economics. For every large message in our network there is at least one small protocol message, so that the absolute increase in the two types may be about equal; however, because there are presently more small than large messages, A is greater than a.

   If we observe that 80% of large messages currently go to or from specialized nodes, and if we believe that a user will generate 50% more messages because of network specialization by year y, then the value for A in the formula for that year should be 1.4. We might plug in values of 1.2, 1.4, and 1.8 to get a range of answers corresponding to a range of assumptions about future network specialization.

2. The effect of intelligent and graphics terminals will be limited almost entirely to the front end of the Network. The use of graphics terminals will increase the large message rate from workers to terminals, because terminal output will sometimes consist of plot information for a full screen instead of one line of text. The effect of intelligent terminals, whether graphics or not, may be complicated. On the one hand, the ability to do local processing, especially screen editing, should result in fewer messages of much larger average size. On the other hand, some users may program their terminals to issue very frequent program or network status checks on a background basis and take some action only when a certain response is obtained, thus greatly increasing the small message rate.

   In any case, the factors b and B representing this effect should probably be multiplicative as are a and A above. Management projections indicate that 1000 of the terminals in the Lab will be intelligent in 10 years. We have observed that, at present, about one-fourth of all terminals are logged in on any morning. An assumption of an upper bound of 2.5 large messages per minute at these terminals gives 625 large messages per minute, which is about half the present rate; thus, we used values of from 1.1 to 1.5 for B. We used values of from 0.9 to 1.1 for b. The small range of values for b indicates that not all terminals will be intelligent and that most messages are already small.

3. The increased use of distributed processors and of local networks will certainly decrease the ICN message rate per user. Almost all of these users' terminal traffic, which consists mostly of small messages, will be eliminated from the ICN. They will still use the ICN for executing large

programs prepared locally and for special services mostly involving large files.

3. Once again, we decided that the factors c and C should be ratios of the present message rates per user. Values of from 0.5 to 1 for C and 0.25 to 1 for c seem reasonable.

4. If electronic mail is implemented using the ICN, then, obviously, message traffic will increase. It is not at all clear that there is any correlation between the rate at which users currently generate messages and the rate at which they will receive mail. However, mail traffic will probably be proportional to the number of people using the system. We assumed the relationship will be linear (although there are certainly other plausible possibilities). Thus we included terms

   $d*Ny$ and $D*Ny$

   in the formulas for the number of small and large messages. We eventually decided that people would send and receive less than five large mailings per day, which is a negligible addition to our load; therefore, we used the value zero for D.

5. The additional message traffic caused by connecting our network to others would depend very much on the administrative nature of the connection. If remote users were given essentially the same capabilities as local users, then the appropriate adjustment to the formulas is simply to increase the value of N by the number of remote users. If use of the connection is restricted to sharing programs, data, and reports between sites, in other words, if the link is used as a fast substitute for the Postal Service, then the message rate might be independent of the number of users altogether and might depend instead on programmatic schedules. We assumed that the latter was more likely and added a simple term e to each formula to account for some small constant number of messages due to this connection.

3.4 Formulas

Collecting all the terms defined in the previous section resulted in the following formulas:

$$SMy = a*b*c*m*Ny + d*Ny + e \tag{1}$$

$$LMy = A*B*C*M*Ny + D*Ny + E \tag{2}$$

where

SMy and LMy	are the number of small and large messages per second in the ICN in year y;
m and M	are the current (1980) number of small and large messages per second per user;
Ny	is the number of CCF users that year;
a	is the factor by which network specialization will affect the number of small messages per second per user that year,
b	represents the effect of intelligent terminals.
c	represents the effect of distributed processing.
d	is the number of small messages per user per second due to electronic mail;
e	is the number of additional small messages per second due to connections to external networks, and
A, B, C, D, and E	are the corresponding factors for large messages

We can now plug various values for each of the factors into the formula and get "best guess," "worst case," and other values for message traffic. We can also experiment with the effects of particular assumptions, for example, we can assume that all terminals will be intelligent in 10 years or that electronic mail traffic is proportional to the square of the number of users. We can investigate "disaster" scenarios; to illustrate, we can determine the rate at which intelligent terminal owners would have to generate status queries to the Network to saturate its message handling capacity. Finally, we can determine by inspection or by trial which assumptions are most critical, for example, the above formulas are clearly more sensitive to the value of a than to the value of e

3.5 Other Possible Factors

The five change factors discussed above re certainly not the only ones that will affect computing in the Laboratory in the next ew years. Since we constructed the above ormulas, we have learned that, unknown to s, others in the Laboratory were already lanning another change, namely a aboratory-wide automated information management system (AIMS). Some of the pieces of uch a system, such as accounting programs nd some inventory programs, are already run n worker computers in the ICN. Their integration into a comprehensive, widely used anagement information system would certainly ncrease network message rates. The point of his example is that as many people as possible, from a variety of disciplines, should be ncluded in the process of thinking of hanges in computing.

More speculative changes than those we ave given might also be included in a projection. Very powerful processors on a single chip will soon be available at very low ost. The use of high-quality graphics output devices may become much more widespread n the Laboratory to display movies (16 rames of graphics output per second) used to tudy simulation modeling programs. Although resent worker computers are not capable of roducing 16 frames per second from these rograms, long sequences of frames could be enerated and stored in CFS; these could be etched and fed to the graphics device by the heap powerful processor at such a rate. If his happens, it will greatly increase the arge message rate.

4. Conclusions

Inserting our "best guess" factor values nto the above formulas resulted in message ates for 1990 of five to six times the resent observed rates. To anyone familiar ith the history of computing, it might seem nlikely that any workload measure on any ystem will grow by "only" 500% in 10 years. f this projection, in fact, turns out to be ow, the reason will probably be that we ailed to anticipate some development in computing that radically affects network use. he necessity of anticipating such changes s, of course, the greatest weakness of our ethod; however, this weakness is inherent to he problem. It can be overcome somewhat by equesting input from as many people as possible.

Our method of prediction presented in his paper identifies specific assumptions. t allows experimenting with different values f factors to see the part each plays in the total prediction. More accurate data about the effect of a given change can be easily incorporated into the formulas so that predictions grow more accurate in an evolutionary way. Concentrating on the effects on individual users might also work well for short-term predictions, but we found this method especially helpful as a way of isolating and organizing the uncertainties and shakey assumptions inherent in long-range prediction.

References

[1] Lindsay, D. S., A Hardware Monitor Study of a CDC KRONOS System, Proceedings, International Symposium on Computer Performance Modeling Measurement and Evaluation, Harvard University, March 1976, pp. 136-144.

[2] Partridge, D. R. and Card, R. E., Hardware Monitoring of Real-Time Computer Systems, Proceedings, International Symposium on Computer Performance Modeling, Measurement and Evaluation, Harvard University, March 1976, pp. 85-101.

[3] McDougall, M. H., The Event Analysis System, Proceedings, 1977 SIGMETRICS/CMG VIII, Washington, D.C., November 1977, pp. 91-100.

[4] Keller, T. W. and Lindsay, D. S., Development of A Performance Evaluation Methodology for the CRAY-1, Proceedings, 1977 SIGMETRICS/CMG VIII, Washington, D.C., November 1977, pp. 169-176.

[5] Abrams, M. D. and Neiman, D. C., NBS Network Measurement Methodology Applied to Synchronous Communications, Proceedings, CPEUG80, Orlando, Florida, November 1980, pp. 63-70.

[6] Hughes, J. H., A Functional Instruction Mix and Some Related Topics, Proceedings, International Symposium on Computer Performance Modeling Measurement and Evaluation, Harvard University, March 1976, pp. 145-153.

[7] Terman, F. W., A Study of Interleaved Memory Systems by Trace Driven Simulation, Proceedings, Symposium on Simulation of Computer Systems, Boulder, Colorado, August 1976, pp. 3-10.

[8] Mamrak, S. A. and Amer, P. D., A Feature Selection Tool for Workload Characterization, Proceedings, 1977 SIGMETRICS/CMG VIII, Washington, D.C., November 1977, pp. 113-120.

[9] Agrawala, A. K. and Mohr, J. M., The Relationship Between the Pattern Recognition Problem and the Workload Characterization Problem, Proceedings, 1977 SIGMETRIC/CMG VIII, Washington, D.C., November 1977, pp. 131-139.

[10] Artis, H. P., A Technique for Establishing Resource Limited Job Class Structure, Proceedings, CMG X, Dallas, Texas, December 1979, pp. 249-253.

[11] Linde, S. and Morgan, L., Workload Forecasting for the Shuttle Mission Simulator Computer Complex at Johnson Space Center, Proceedings, CMG XI, Boston, Massachusetts, December 1980, pp. 10-15.

[12] Mohr, J. and Penansky, S. G., A Framework for the Projection of ADP Workloads, Proceedings, CMG XI, Boston, Massachusetts, December 1980, pp. 5-9.

[13] Perlman, W. and Jozwik, R., DP Workload Forecasting, Proceedings, CMG X, Dallas, Texas, December 1979, pp. 469-475.

[14] Smith, C. and Browne, J. C., Modeling Software Systems for Performance Predictions, Proceedings, CMG X, Dallas, Texas, December 1979, pp. 321-342.

[15] Smith, C. and Browne, J. C., Performance Specifications and Analysis of Software Designs, Proceedings, Conference on Simulation Measurement and Modeling of Computer Systems, Boulder, Colorado, August 1979, pp. 173-182.

[16] Axelrod, C. W., The Dynamics of Computer Capacity Planning, CMG Transactions, Number 29, September 1980, pp. 3.3-3.21.

[17] Tomberlin, R. D., Forecasting Computer Processing Requirements: A Case Study, Proceedings, CPEUG80, Orlando, Florida, November 1980, pp. 255-261.

[18] Blood, M., Christman, R., and Collins, B., Experience with the LASL Common File System, Digest of Papers from Fourth IEEE Symposium on Mass Storage Systems, Denver, Colorado, April 1980, pp. 51-54.